arXiv:1701.00505v2 [stat.ME] 13 Jan 2017

Statistical Inference for Network Samples Using Subgraph Counts

P-A. G. Maugis[1], C. E. Priebe[2], S. C. Olhede[1], and P. J. Wolfe[1]

[1] *University College London – Department of Statistical Science*
[2] *Johns Hopkins University – Department of Applied Mathematics and Statistics*

Abstract

We consider that a network is an observation, and a collection of observed networks forms a sample. In this setting, we provide methods to test whether all observations in a network sample are drawn from a specified model. We achieve this by deriving, under the null of the graphon model, the joint asymptotic properties of average subgraph counts as the number of observed networks increases but the number of nodes in each network remains finite. In doing so, we do not require that each observed network contains the same number of nodes, or is drawn from the same distribution. Our results yield joint confidence regions for subgraph counts, and therefore methods for testing whether the observations in a network sample are drawn from: a specified distribution, a specified model, or from the same model as another network sample. We present simulation experiments and an illustrative example on a sample of brain networks where we find that highly creative individuals' brains present significantly more short cycles.

KEYWORDS: Statistical Testing, Graphon Model, Subgraph Statistics.

This work was supported in part by DARPA SIMPLEX contract N66001-15-C-4041 the US Army Research Office under Multidisciplinary University Research Initiative Award 58153-MA-MUR; by the US Office of Naval Research under Award N00014-14-1-0819; by the UK Engineering and Physical Sciences Research Council under Mathematical Sciences Leadership Fellowship EP/I005250/1, Established Career Fellowship EP/K005413/1, Developing Leaders Award EP/L001519/1, and Award EP/N007336/1; by the UK Royal Society under a Wolfson Research Merit Award; and by Marie Curie FP7 Integration Grant PCIG12-GA-2012-334622 and the European Research Council under Grant CoG 2015-682172NETS, both within the Seventh European Union Framework Program. The authors thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, UK, for support and hospitality during the program Theoretical Foundations for Statistical Network Analysis (EPSRC grant no. EP/K032208/1) where a portion of the work on this paper was undertaken. The authors thank Joshua T. Vogelstein and his team for connectome data and expertise.

1 Introduction

We show that subgraph counts are flexible and powerful statistics for inference on collections of networks. Specifically, we use subgraph counts to test the hypotheses that all networks in a sample are generated either from a given distribution, from distributions in a given model, or from the same model as that of another sample.

Our results address the inference problem raised by the following experiment [1]: The networks connecting brain regions of individuals of varied levels of creativity is observed. However, while observations can be assumed to be independent, due to the variability of the brain structure and the instability of the observation technique, they cannot be assumed to be identically distributed; for instance, they need not contain as many nodes and edges. How, while allowing for such variations, can we test for significant differences between individuals with different levels of creativity?

Formally, we consider that a network is an observation—say G_i—and a collection of observed networks form a sample—say $\mathcal{G} = (G_1, \ldots, G_N)$. Then, our goal is to infer distributional properties of the G_i-s as N grows. This parallels more classical statistical settings, where an observation is a vector—such as $X_i \in \mathbb{R}^k$—and a sample is a matrix: $\mathcal{X} = (X_1, \ldots, X_N) \in \mathbb{R}^{k \times N}$. However, our setting strongly differs from the one where only one very large network is observed, and for which many methods already exist (see [2–10], to cite but a few). Surprisingly, no statistical method exists to compare samples of small networks, and currently only tools to compare two large networks are available [11, 12].

Here we provide an analog of a multivariate t-test for network samples: methods to test whether a given network sample $\mathcal{G}$ presents *averages* consistent with either a specific model, or with that of another sample. The averages we use are *subgraph counts*; e.g., the number of or in the sample. The choice of subgraph counts as statistics is motivated by their success in comparing large networks [13, 14], but also by results in random graph theory and the study of large graphs. In both fields, subgraph counts have proved to be the most powerful tool available to compare networks [15, 16], and are known to have properties similar to moments of random variables [17].

Formally, to perform our tests, we are first embedding network samples into a space defined by subgraph counts. While related testing procedures also use embeddings [11, 12], using subgraph counts presents three key advantages: first, if the G_i-s are generated by a blockmodel [2]—the most popular random network model to date—and for an appropriate family of subgraphs, the embedding is one-to-one. This result is known as the finite forciblity of blockmodels [17]. Second, very few assumptions on each G_i need to be made as N grows to obtain consistency and asymptotic normality of the image of $\mathcal{G}$ in the embedding space. This enables us to work under a very flexible null model. Finally, because it relies on physical properties of the G_i-s—the number of , , and so on—this embedding remains interpretable.

In the remainder of this article, we first introduce subgraph counts and the graphon model. We then present, successively, the case where all the networks

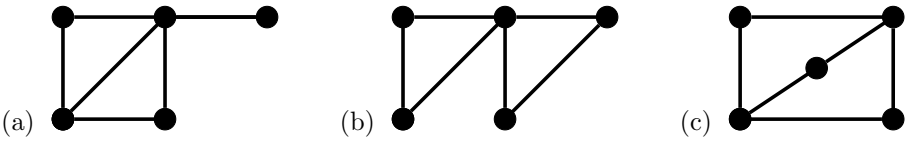

Figure 1: Example of subgraph counts. There are 6 copies of the edge (⊷) in all three graphs. There are 2 copies of the triangle (△) in a) and b), but 0 in c). There are 1, 0 and 3 copies of the square (□) in a), b) and c) respectively.

in the sample come from the same graphon model (but are not necessarily of the same size), and the case where each observed network may come from a different graphon. In both cases, we prove asymptotic normality of our estimator, and present representative examples showing the practical use of the result. We conclude with an application to connectomes, and a discussion.

2 Subgraph Counts in the Graphon Model

We now define our statistics (subgraph counts) and our null model (the graphon model). Subgraph counts are natural statistics to compare networks for two reasons. First, subgraph counts intuitively summarize a network through its fundamental building blocks. This has historically given them purchase to address hard fundamental and empirical problems [13, 14, 18]. Second, subgraph counts present tractable analytical properties. We will describe and leverage these properties below, in a manner paralleling what is done in related literatures [4, 6, 18].

A *subgraph count* is the number of copies of a given graph in another graph (see Fig. 1). Throughout, we call the subgraph—and denote F—the graph which is counted and G the larger graph in which the counting takes place. All graphs will be simple (unweighted, no self loops or multiple edges). Subgraphs are also termed motifs, pattern graphs or shapes depending on the field [13, 19–21].

For clarity, we define subgraph counts formally as follows (we write $|F|$ to be the number of nodes in F, and $F \subset G$ if F is a (not necessarily induced) subgraph of G):

Definition 1 (Subgraph count $X_F(G)$)**.** *For two graphs F and G, we call the count of F in G—and write $X_F(G)$—the following quantity:*

$$X_F(G) = \#\big\{F' \subset G \,:\, F' \text{ is a copy of } F\big\},$$

where F' is a copy of F if there exists an adjacency preserving bijection between F and F'. Furthermore, for a tuple $\mathcal{F} = (F_1, \ldots, F_k)$ of subgraphs, we write $X_{\mathcal{F}}(G)$ to be the vector $(X_{F_1}(G), \ldots, X_{F_k}(G))$.

With this notation, calling G_a, G_b and G_c the graphs in Fig. 1, we have that $X_{\square}(G_a) = 1$, $X_{\square}(G_b) = 0$, $X_{\square}(G_c) = 3$.

The power of subgraph counts to study networks stems from their inherent linearity. Indeed, products of subgraph counts are but linear combinations of other subgraph counts. Intuitively, a product of two subgraph counts will involve counting pairs of copies, and therefore can be recovered by counting the number of copies of all subgraphs that can be induced by a pair of copies. More precisely, in the Appendix we show the following:

Lemma 1 (Linearity of subgraph counts)**.** *For any two graphs F and F', there are factors c_H and a set $\mathcal{H}_{FF'}$ of subgraphs—the set of subgraphs that can be obtained using one copy of each F and F' as building blocks—such that for any graph G*

$$X_{F_1}(G)X_{F_2}(G) = \sum_{H \in \mathcal{H}_{FF'}} c_H X_H(G).$$

For instance, as these will be used later on, we have that in *any* graph G,

$$X_{\triangle}(G)^2 = 2X_{\triangle\triangle}(G) + 2X_{\wedge\wedge}(G) + 2X_{\square}(G) + X_{\triangle}(G),$$
$$X_{\square}(G)^2 = 2X_{\square\square}(G) + 2X_{\square\triangle}(G) + 6X_{\diamond}(G) + 2X_{\square}(G)$$
$$+ 2X_{\square\square}(G) + 2X_{\triangle\triangle}(G) + 6X_{\triangle}(G) + 6X_{\square}(G) + X_{\square}(G).$$

This algebraic property of subgraph counts underpins the proofs of [6, 16, 18], and the subgraph counting algorithms of [20, 21], among many other examples. Crucially, as opposed to cases where the model enforces linearity—such as with assumptions of Normality—it is the nature of the statistics (subgraph counts) and the system (graphs) that makes the problem linear.

The linearity of subgraph counts allows us to use as null the very flexible graphon model [17]. This framework subsumes most models used in the statistical literature on networks; e..g, blockmodel [2] and dot-product models [5]. It has the intuitive structure of affixing to each node i a latent feature (here x_i) and of connecting nodes i and j (conditionally independently) with a probability determined by the node features (here $f(x_i, x_j)$).

Definition 2 (Graphon f and random graph $G_n(f)$)**.** *Fix a symmetric integrable map $f : [0,1]^2 \to [0,1]$, and call it a graphon. We call $G_n(f)$ the random graph distribution over graphs with n nodes such that: to each node is randomly and independently assigned a feature $x_i \in [0,1]$, with $x_i \sim \text{Unif}([0,1])$; and where edges form independently conditionally on $\{x_i\}_{i\in[n]}$ with probability*

$$\mathbb{P}[ij \in G | x_i, x_j] = f(x_i, x_j).$$

To recover a blockmodel with K blocks, it suffices to consider a partition of $[0,1]$ in K sets (i.e., $(P_1, \ldots, P_K) \in \mathcal{P}_K([0,1])$) and set f as constant over each $P_u \times P_v$. The dot-product model is recovered with a graphon f of finite rank; i.e., $f(x,y) = \sum_{u \le K} \lambda_u f_u(x) f_u(y)$.

In the graphon framework, subgraph counts have direct interpretation as moments of f [17]. Specifically, if $G \sim G_n(f)$, then the moments of $X_F(G)$ are

moments of f. Therefore, following results similar to the Hausdorff moment problem, subgraph counts are sufficient statistics to distinguish between any two graphons [16, 17]. However, there are no guarantees on which subgraphs are needed to distinguish between two graphons. For blockmodels and finite rank models, we know only that a finite number is sufficient (a concept know as finite forcibility, see [17, Chapter 16.7 & Appendix 4] for more details).

Unfortunately, all known results on subgraph counts under the graphon model consider the setting where one very large graph is observed. Here we present the tools to address the problem where a sample of small graphs is observed.

3 The simple case: Samples from one graphon

We now present a central limit theorem as well as practical methods to build confidence regions for the subgraph counts observed in a network sample $\mathcal{G} = (G_1, \ldots, G_N)$. In this section, we assume that there is a graphon f such that each G_i is drawn independently from $G_{n_i}(f)$ (where $n_i = |G_i|$).

Fix $F \in \mathcal{F}$ and $G \in \mathcal{G}$. In this setting, $X_F(G)$ is a random variable, and the first parameter to consider is its mean. To compute this mean, let $F_1, \ldots F_m$ be all the copies of F in K_G (the complete graph over the nodes of G), so that using the linearity of the expectation, we have that

$$\mathbb{E}X_F(G) = \mathbb{E}\sum_{j\in[m]} 1_{\{F_j \subset G\}} = \sum_{j\in[m]} \mathbb{E}1_{\{F_j\subset G\}}.$$

Then, direct computations show that $\mathbb{E}1_{\{F_j\subset G\}}$ does not depend on j (see Proposition A.1), and that

$$\mathbb{E}1_{\{F_j\subset G\}} = \mu_F(f) := \int_{[0,1]^{|F|}} \prod_{uv\in F} f(x_u, x_v) \prod_{u\in F} dx_u,$$

so that $\mathbb{E}X_F(G) = m\mu_F(f)$. Observe that $\mu_F(f)$ is a moment of the graphon f, as discussed above.

Similar computations for higher moments, aided by Lemma 1, enable us to use the Lindeberg-Feller central limit theorem along with the Cramer-Wold device to obtain the following:

Theorem 1 (Statistical properties of subgraph counts)**.** *Fix a tuple of graphs $\mathcal{F}$, a graphon f and a sequence $n = (n_i)_{i\in\mathbb{N}}$ such that $2\max_{F\in\mathcal{F}}|F| \leq \min_{i\in\mathbb{N}} n_i$. Let $\mathcal{G} = (G_i)_{i\in[N]}$ be a graph sample such that for all i, $G_i \sim G_{n_i}(f)$. Set $\hat{\mu}_F(\mathcal{G}) = N^{-1}\sum_{G\in\mathcal{G}} X_F(G)/X_F(K_G)$, $\hat{\mu}_{\mathcal{F}}(\mathcal{G}) = (\hat{\mu}_F(\mathcal{G}))_{F\in\mathcal{F}}$, and $\mu_{\mathcal{F}}(f) = (\mu_F(f))_{F\in\mathcal{F}}$. Then, $\hat{\mu}_{\mathcal{F}}(\mathcal{G})$ is an unbiased, $\sqrt{N}$-consistent and asymptotically normal estimator of $\mu_{\mathcal{F}}(f)$; i.e., $\mathbb{E}\hat{\mu}_{\mathcal{F}}(\mathcal{G}) = \mu_{\mathcal{F}}(f)$ and there exists $\Sigma_{\mathcal{F}}(n, f)$ such that asymptotically in N:*

$$\sqrt{N}\big(\hat{\mu}_{\mathcal{F}}(\mathcal{G}) - \mu_{\mathcal{F}}(f)\big) \to \text{Normal}\big(0, \Sigma_{\mathcal{F}}(n, f)\big).$$

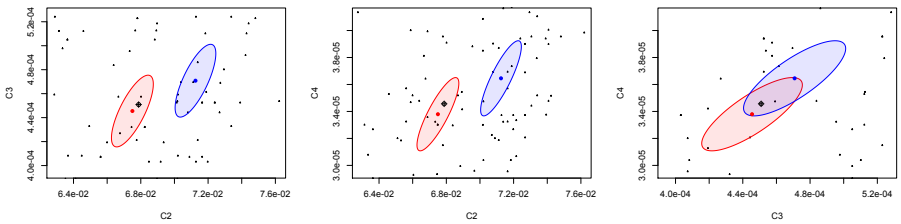

Figure 2: Testing for a blockmodel using $\mathcal{F} = \{C_2, C_3, C_4\} = \{\bullet\!\!-\!\!\bullet, \triangle, \square\}$. The sample $\mathcal{G}$ is such that: $N = 300$, n_i is the i-th digit of π plus 40, G_i is drawn from a two block graphon f ($G_i \sim G_{n_i}(f)$). The density estimate $\hat{\mu}_{\mathcal{F}}(\mathcal{G})$ is denoted by a black cross. Overlaid are the expected densities (colored dots) and the confidence ellipse (shaded area) for two alternative graphons f_a and f_b. The p-values obtained with the Mahalanobis distance are respectively 0.6 and 7e−15.

Furthermore, for each $F, F' \in \mathcal{F}$,

$$\operatorname{cov}(\hat{\mu}_F(\mathcal{G}), \hat{\mu}_{F'}(\mathcal{G})) = \sum_{H \in \mathcal{H}_{FF'} \setminus \{F \sqcup F'\}} \omega_H(n; N)\big(\mu_H(f) - \mu_F(f)\mu_{F'}(f)\big), \quad (1)$$

with $F \sqcup F'$ *the disjoint union of* F, F', *and* $\omega_H(n; N) = \frac{1}{N} \sum_{i=1}^{N} \frac{c_H X_H(K_{n_i})}{X_F(K_{n_i}) X_{F'}(K_{n_i})}$.

Crucial to the following is the covariance matrix $\Sigma_{\mathcal{F}}(n, f)$—which can be obtained by taking the limit in N in (1) for each $F, F' \in \mathcal{F}$—and which will enable the computation of confidence regions. Interestingly, its elicitation is more involved than for the study of large graphs, where only a few terms dominate. We refer to the Appendix for the proof as well as a simulation experiment.

Theorem 1 enables testing against the null that all G_i are drawn from a given graphon. To make this concrete, we consider an example in Fig. 2. There, we observe a graph sample $\mathcal{G} = (G_1, \ldots, G_{300})$, and aim to compare it to two graphon models f_a (in red) and f_b (in blue) using Theorem 1; i.e., we assume that for $i \in [N]$, $G_i \sim G_{n_i}(f)$ and consider the null hypothesis $H_0 : f = f_a$ and the alternative $H_1 : f = f_b$. We draw as a black cross $\hat{\mu}_{\mathcal{F}}(\mathcal{G})$ and as smaller black points the $\mu_{\mathcal{F}}(G_i)$. The sizes of the networks in $\mathcal{G}$, the n_i, are non-random but not constant. We achieve this by using the sequence of digits of π.

First, since we have specified f_a and f_b, we can evaluate both $\mu_{\mathcal{F}}(f_a)$ and $\mu_{\mathcal{F}}(f_b)$ and draw them on the figure (as a red and blue dot respectively). Then, since $n = (n_i)_{i \leq N}$ is observed, we can compute $\Sigma_{\mathcal{F}}(n, f_a)$ and $\Sigma_{\mathcal{F}}(n, f_b)$ using Theorem 1, which allows us to compute the confidence ellipse around $\mu_{\mathcal{F}}(f_a)$ and $\mu_{\mathcal{F}}(f_b)$ (in shaded red and blue respectively). Finally, since we know the limit distribution and covariance under the null, we can use $\Sigma_{\mathcal{F}}(n, f_a)$ and $\mu_{\mathcal{F}}(f_a)$ to compute a p-value using Mahalanobis distance.

In the following we consider the case where instead of testing against the null of a single graphon, we test against the null of a graphon class.

4 The general case: Flexible sampling design

Here we expand our results to cases where the observed networks may be generated from different graphons. Indeed, in many settings, the sampling mechanism may distort the structure of the underlying graphon; e.g., although the network connecting brain regions can be satisfactorily modeled by a blockmodel [22], the proportion of nodes of each block may be different in different experimental settings, so that each observation is drawn from a different blockmodel.

In this practically important and conceptually challenging new setting, the proof techniques developed for Theorem 1 yield the following.

Theorem 2. *Fix a tuple of graphs $\mathcal{F}$, a sequence of graphons $f = (f_i)_{i\in\mathbb{N}}$ and a sequence of integers $n = (n_i)_{i\in\mathbb{N}}$ such that $2\max_{F\in\mathcal{F}}|F| \leq \min_{i\in\mathbb{N}} n_i$. Let $\mathcal{G} = (G_i, \dots, G_N)$ be a graph sample such that for all i, $G_i \sim G_{n_i}(f_i)$. Set $\hat{\mu}_{\mathcal{F}}(\mathcal{G}) = N^{-1}\sum_{G\in\mathcal{G}}(X_F(G)/X_F(K_G))_{F\in\mathcal{F}}$ and $\mu_{\mathcal{F}}(f;N) = N^{-1}\sum_{i=1}^{N}\mu_{\mathcal{F}}(f_i)$. Then, asymptotically in N, and for some matrix $\Sigma^*_{\mathcal{F}}(n,f)$, we have that*

$$\sqrt{N}\big(\hat{\mu}_{\mathcal{F}}(\mathcal{G}) - \mu_{\mathcal{F}}(f;N)\big) \to \mathrm{Normal}\big(0, \Sigma^*_{\mathcal{F}}(n,f)\big).$$

Therefore, even in this much more flexible setting, we can recover the barycenter of the $\mu_{\mathcal{F}}(f_i)$. However, the variance has now a more complex structure, and we refer to the Appendix for details.

Following the intuition of our example of brain networks, and to make the usefulness of Theorem 2 concrete, we introduce the *flexible stochastic blockmodel* (FSBm).

Definition 3 (FSBm and embedding shape)**.** *For a symmetric matrix $B \in [0,1]^{K\times K}$ we call $\mathcal{D}(B)$ the set of all possible graphons with the same block structure as B; i.e.,*

$$\mathcal{D}(B) = \{f : \exists(P_1,\dots,P_K) \in \mathcal{P}_K([0,1])\, s.t. \forall x \in P_s, y \in P_t, f(x,y) = B_{st}\}.$$

For a tuple $\mathcal{F}$ of graphs, we call embedding shape the set

$$\mu_{\mathcal{F}}(B) = \{\mu_{\mathcal{F}}(f) \text{ for } f \in \mathcal{D}(B)\}.$$

For instance, with $B \in [0,1]^{2\times 2}$ and $\mathcal{F} = \{\bullet\!\!-\!\!\bullet, \wedge\}$, then:

$$\mu_{\mathcal{F}}(B) = \big\{\big(\pi^2 B_{11} + 2\pi(1-\pi)B_{12} + (1-\pi)^2 B_{22},\\ \pi^3 B_{11}^3 + 3\pi^2(1-\pi)B_{11}B_{12}^2 + 3\pi(1-\pi)^2 B_{22}B_{12}^2 + (1-\pi)^3 B_{22}^3\big)\\ : \pi \in [0,1]\big\}.$$

The most direct way of using the FSBm is to test for all the f_i being equal to any blockmodel instance in a class; i.e assume that all G_i-s are drawn from a graphon f and test for the null $H_0 : f \in \mathcal{D}(B)$. This is achieved by using a composite hypothesis test, and our results allow us to produce confidence regions and p-values using the same tools as before.

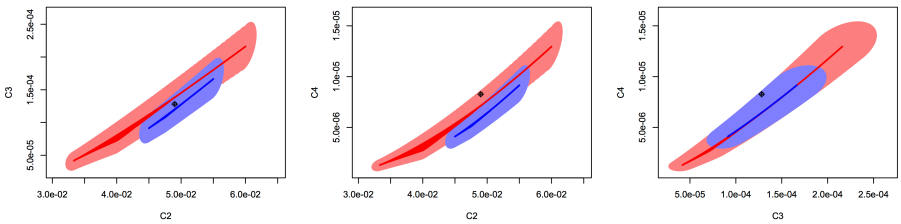

Figure 3: Testing for a FSBm class. The sample $\mathcal{G}$ is such that: $N = 200$, n_i is the i-th digit of π plus 30, G_i is drawn from a graphon f ($G_i \sim G_{n_i}(f)$). With $\mathcal{F} = \{C_2, C_3, C_4\}$, we estimate $\hat{\mu}_{\mathcal{F}}(\mathcal{G})$, and plot it as a black cross. Then, we draw in solid color the embedding shapes $\mu_{\mathcal{F}}(B_a)$ and $\mu_{\mathcal{F}}(B_b)$. In shaded color we draw the associated confidence regions; p-values can be obtained using the Mahalanobis distance associated with the closest point to $\hat{\mu}_{\mathcal{F}}(\mathcal{G})$ in $\mu_{\mathcal{F}}(B_a)$ and $\mu_{\mathcal{F}}(B_b)$.

We present such an example in Fig. 3. There, we observe $\mathcal{G} = (G_1, \dots, G_{200})$, and consider two FSBm classes generated from B_a (in red) and B_b (in blue). Then, we assume that all networks in the sample are drawn from a graphon f and test for the null $H_0 : f \in \mathcal{D}(B_a)$ and the alternative $H_1 : f \in \mathcal{D}(B_b)$. We first represent $\mu_{\mathcal{F}}(\mathcal{G})$ as a black cross. Using Definition 3, we plot the embedding shapes $\mu_{\mathcal{F}}(B_a)$ and $\mu_{\mathcal{F}}(B_b)$ in solid red and blue respectively. The confidence regions (in shaded red and blue) are the union of the confidence ellipses at all points in $\mu_{\mathcal{F}}(B_a)$ and $\mu_{\mathcal{F}}(B_b)$.

A more general use of Theorem 2 is to test for all graphs in a sample being drawn from elements of a FSBm class; i.e assume that the G_i-s are drawn from the f_i-s and test for the null $H_0 : \forall i \in [N], f_i \in \mathcal{D}(B)$ for some B. As before, we face a composite null, and we may compute the confidence region and the p-value by scanning all possible sequences f. This, however, is clearly computationally intractable. Nonetheless, the form of the variance and the structure of the FSBm allows us to propose conservative confidence regions and p-values that can be efficiently computed (we fully describe the method in the Appendix).

We present an example in Fig. 4. There we observe $\mathcal{G} = (G_1, \dots, G_{10^3})$ and consider two FSBm classes generated by B_a and B_b. We first plot $\mu_{\mathcal{F}}(\mathcal{G})$ as a black cross. Then, using Definition 3, we plot the convex hull of the embedding shapes $\mu_{\mathcal{F}}(B_a)$ and $\mu_{\mathcal{F}}(B_b)$ (in solid red and blue respectively) wherein—by Theorem 2—$\mu_{\mathcal{F}}(f; N)$ must lie. Finally, we use a method described in the Appendix to produce the confidence region around each shape (in shaded color).

5 Application: Are creative brains different?

We now consider a sample of brain networks $\mathcal{G} = (G_1, \dots, G_{114})$ [23]. This sample was produced in two steps: first, magnetic resonance images of each 113 subjects'

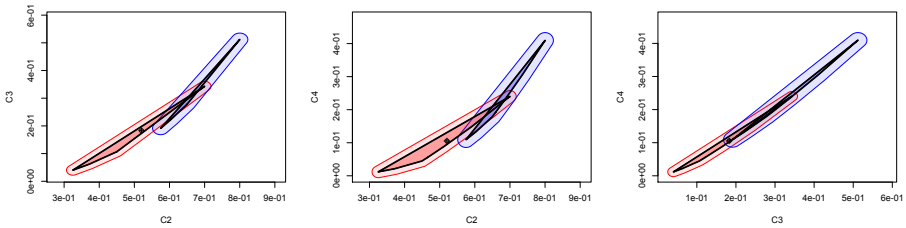

Figure 4: Testing for a full FSBm class. The sample $\mathcal{G}$ is such that: $N = 10^3$, n_i is the i-th digit of π plus 50, G_i is drawn from a graphon f_i ($G_i \sim G_{n_i}(f_i)$). With $\mathcal{F} = \{C_2, C_3, C_4\}$, we estimate $\hat{\mu}_{\mathcal{F}}(\mathcal{G})$, and plot it as a black cross. Then, we draw in solid color the convex hulls of the embedding shapes $\mu_{\mathcal{F}}(B_a)$ and $\mu_{\mathcal{F}}(B_b)$. In shaded color we draw the associated confidence regions; approximate (and conservative) p-values can be obtained by determining the confidence level at which the observation ceases to be in the confidence region.

brains were taken; then, the networks connecting each subjects' brain regions were estimated using these images [1]. All networks in the sample contain 70 nodes. Furthermore, we have available a covariate $C = (c_1, \ldots, c_{113})$ measuring the subjects' creativity.

To study this network sample and use the covariate C, we introduce a direct extension of our results to compare two network samples:

Corollary 1 (Two-sample test)**.** *Fix a tuple of subgraphs $\mathcal{F}$ and two network samples $\mathcal{G}$ and $\mathcal{G}'$ generated respectively by the graphons f and f' and the network size sequences n and n'. Then, as both $|\mathcal{G}|$ and $|\mathcal{G}'|$ tend to infinity, and if $\min(n, n') \geq 2 \max_{F \in' F} |F|$, we have that if $f = f'$, then*

$$\frac{\sqrt{|\mathcal{G}||\mathcal{G}'|}}{\sqrt{|\mathcal{G}| + |\mathcal{G}'|}} \big(\hat{\mu}_{\mathcal{F}}(\mathcal{G}) - \hat{\mu}_{\mathcal{F}}(\mathcal{G}')\big) \to \text{Normal}\big(0, \Sigma_{\mathcal{F}}(n, n', f)\big),$$

and $\Sigma_{\mathcal{F}}(n, n', f)$ may be estimated at rate $\sqrt{|\mathcal{G}| + |\mathcal{G}'|}$ from the samples.

Unfortunately, estimating $\Sigma_{\mathcal{F}}(n, n', f)$ requires counting subgraphs of order $2 \max_{F \in \mathcal{F}} |F|$, making the procedure computationally intensive. This compels us to work with $\mathcal{F} \subset \{$ $\}$. Furthermore, although our estimator of $\Sigma_{\mathcal{F}}(n, n', f)$ is entrywise Normal, it is not Wishart as in the classical setting. Therefore, we have no guarantee of the estimate being positive definite, and cannot use the Hotelling's T-squared distribution to compute p-values. If the estimate is positive definite, we recommend ignoring the variations in $\Sigma_{\mathcal{F}}(n, n', f)$ and using the $\chi^2_{|\mathcal{F}|}$ distribution. If the estimate fails to be positive definite, we recommend using only the marginals.

Before analyzing $\mathcal{G}$ using our results, we make the following test: we subsample uniformly at random and without replacement from $\mathcal{G}$, yielding $\mathcal{G}_1$ and $\mathcal{G}_2$ such that $\mathcal{G}_1 \cup \mathcal{G}_2 = \mathcal{G}$, and use Corollary 1 to test for $\mathcal{G}_1$ and $\mathcal{G}_2$ being drawn from the

Quantile (q)	p-value		
0.5	0.126	0.110	0.115
0.4	0.077	0.051	0.050
0.3	0.062	0.042	0.040
0.2	0.014	0.011	0.012
0.1	0.046	0.047	0.061

Table 1: Testing for differences between $\mathcal{G}_1^q$ and $\mathcal{G}_2^q$. For each $F \in \{$, , $\}$ and $q \in \{0.1, 0.2, \ldots, 0.5\}$ we produce the p-value for the null $H_0 : \mu_F(\mathcal{G}_1^q) = \mu_F(\mathcal{G}_2^q)$. The p-values increase with q, except for $q = 0.1$, in which case $|\mathcal{G}_1^q|$ and $|\mathcal{G}_2^q|$ are too small for the test to be significant.

same graphon f. Unless $\mathcal{G}$ presents characteristics that cannot be explained by our results, $\mathcal{G}_1$ and $\mathcal{G}_2$ should be indistinguishable, and we expect to see p-values that are uniformly distributed in $[0, 1]$.

We perform this experiment 100 times, and obtain a sample of p-values for which we fail to reject the null of a uniform distribution using the Kolmogorov–Smirnov test ($D = 0.09$, p-value $= 0.3$). For this test we use $\mathcal{F} = \{$$\}$ because of the small sample ($|\mathcal{G}_1| + |\mathcal{G}_2| = 113$) size and a very high level of correlation; otherwise the estimated covariance matrix often failed to be positive definite.

We now use C to split $\mathcal{G}$ in two samples. To do so, we choose to build a first subsample $\mathcal{G}_1$ containing the less creative, and a second subsample $\mathcal{G}_2$ containing the more creative. More precisely, for a quantile q and denoting Q_C the empirical quantile function of C:

$$\mathcal{G}_1^q = \{G_i \in \mathcal{G} \,:\, c_i \leq Q_C(q)\} \quad \& \quad \mathcal{G}_2^q = \{G_i \in \mathcal{G} \,:\, c_i > Q_C(1-q)\}.$$

Interestingly, for $q = 0.5$ and $q = 0.4$, we fail to reject the null that the networks in $\mathcal{G}_1^q$ and $\mathcal{G}_2^q$ come from the same graphon (see Table 1 for the p-values). However, for $q = 0.3$ we can reject the null of the same graphon at the 5% confidence level using or , but not .

Thus, we observe that individuals with a very high level of creativity present significantly more and than those with a very low level of creativity. We now aim to understand whether the added and arise from a few edges completing partially present shapes or from fully new and . To do so, we first observe that if $G \sim G_n(f)$, then $\overline{G} \sim G_n(1-f)$, where $\bar{G}$ is the complement graph of G. Therefore, we may use our tests on $\overline{G}$, which can be understood as estimating $\mu_{\mathcal{F}}(1-f)$ instead of $\mu_{\mathcal{F}}(f)$ to compare network samples.

Then, using the $\overline{\mathcal{G}}_i^q = \{\overline{G} \,:\, G \in \mathcal{G}_i^q\}$, we can test whether there are significantly more fully absent subgraphs in $\mathcal{G}_1^q$ compared to $\mathcal{G}_2^q$. There, we find we cannot reject this null; i.e., we cannot reject the null of the networks in $\overline{\mathcal{G}}_1^q$ and $\overline{\mathcal{G}}_2^q$ coming from the same graphon for $q \geq 0.3$. Therefore, we conclude that the added and in the highly creative arise from a few edges completing partially present and .

6 Discussion

We provide the tools to perform statistical inference on a network sample using subgraph counts. Our two main results provide consistency and asymptotic normality of subgraph counts under very flexible conditions. Using these results, we show that subgraph counts are powerful statistics to test whether network samples come from a specified distribution, a specified model, or from the same model.

The key insight we provide is that statistical inference methods paralleling classical ones for standard samples may be obtained for network samples. From this perspective, our results may be seen as providing an analog of a multivariate t-test for network samples. However, going beyond what our results directly imply, we expect that parallels to ANOVA, model selection, model ranking, and goodness of fit may be obtained for network samples using our proof techniques.

Unfortunately, practical implementation of Corollary 1, and of any other expansion on our results, would present hard computational problems. In effect, a fully automatic algorithm to estimate $\Sigma_{\mathcal{F}}(f)$ would need to first automatically build the set $\mathcal{H}_{FF'}$ and then compute the $X_H(G)$ for each H in that set. Such an algorithm would require a number of operations growing combinatorially in the number of nodes in F and F' and exponentially in the total number of nodes in the sample. Thus, even if each graph in the sample is small, such methods would face hard-to-solve computational bottlenecks. Using state of the art subgraph counting methods available today allows consideration of subgraphs F and F' over at most four nodes [21, 24].

It follows that, although we break the statistical inference bottleneck regarding network samples, the computational challenges now become only more salient. At this point, whether there exists compromises allowing for both rigorous statistical inference and scalable computation is the fundamental open question.

APPENDIX A: Properties of subgraph counts

In the following we formalize certain notions we use loosely in the main body (especially the notion of copy and the sets $\mathcal{H}_{FF'}$, as well as the constants c_H), prove all our results, and provide more details on the numerical examples we present.

We start by formally introducing the notion of graph equivalence.

Definition A.1 (Graph equivalence '$\equiv$')**.** *Fix two graphs F and F'. We say that F is equivalent to—or is a copy of—F', and write $F \equiv F'$, if there exists a bijective map ϕ from the vertex set of F to the vertex set of F' such that ij is an edge in F if and only if $\phi(i)\phi(j)$ is an edge in F'.*

We can now provide a more complete, but otherwise equivalent, definition of number of copies.

Definition A.2. *Fix two graphs F and G. We denote $X_F(G)$ the number of non-necessarily induced subgraphs of G equivalent to F; i.e,*

$$X_F(G) = \#\left\{F' \subset G : F' \equiv F\right\},$$

where $F \subset G$ if both the vertex end edge sets of F are subsets of the vertex and edge sets of G respectively.

With this notation, we may prove our first lemma, establishing the linearity of subgraph counts. We first define $\mathcal{H}_{F_1F_2}$ and c_H, generalizing definitions given in [18].

Definition A.3 (Overlapping copies)**.** *For two graphs F_1 and F_2 we write $\mathcal{H}_{F_1F_2}$ the set of unlabeled graphs that can be formed by two copies of F_1 and F_2, and c_H the number of ways a given $H \in \mathcal{H}_{F_1F_2}$ can be built from copies of F_1 and F_2:*

$$\begin{cases} \mathcal{H}_{F_1F_2} & = \left\{ H \subset K_{|F_1|+|F_2|} \ : \quad \begin{array}{c} \exists F_1', F_2' \subset K_{|F_1|+|F_2|} \text{ s.t.,} \\ F_1' \equiv F_1, F_2' \equiv F_2, \\ \text{and } H = F_1' \cup F_2' \end{array} \right\} / \equiv \\ c_H & = \#\left\{ (F_1', F_2') \subset H \ : \quad \begin{array}{c} F_1' \equiv F_1, F_2' \equiv F_2 \\ \text{and } H = F_1' \cup F_2' \end{array} \right\}. \end{cases}$$

Finally, call $\mathcal{H}^*_{F_1F_2}$ the set $\mathcal{H}_{F_1F_2}$ removed of $F_1 \sqcup F_2$, the vertex disjoint union of F_1 and F_2.

Lemma A.1 (Copies pairwise interaction)**.** *Fix three graphs F_1, F_2 and G. Then,*

$$X_{F_1}(G)X_{F_2}(G) = \sum_{H \in \mathcal{H}_{F_1F_2}} c_H X_H(G).$$

Proof. We start by writing:

$$\begin{aligned} X_{F_1}(G)X_{F_2}(G) &= \sum_{F_1' \subset G} 1_{\{F_1' \equiv F_1\}} \sum_{F_2' \subset G} 1_{\{F_2' \equiv F_2\}} \\ &= \sum_{\substack{F_1' \subset G \\ F_2' \subset G}} 1_{\{F_1' \equiv F_1\}} 1_{\{F_2' \equiv F_2\}}. \end{aligned} \tag{2}$$

Now, from Definition A.3, we first note that by construction of $\mathcal{H}_{F_1F_2}$, for each pair F_1', F_2' in the sum, $1_{\{F_1' \equiv F_1\}} 1_{\{F_2' \equiv F_2\}} = 1$ if and only if there exists $H \in \mathcal{H}_{F_1F_2}$ such that $F_1' \cup F_2' \equiv H$. Therefore we can reindex the sum in (2) as follows:

$$\begin{aligned} X_{F_1}(G)X_{F_2}(G) &= \sum_{F_1', F_2' \subset G} 1_{\{\exists H \in \mathcal{H}_{F_1F_2} \,:\, F_1' \cup F_2' \equiv H\}} \\ &= \sum_{H \in \mathcal{H}_{F_1F_2}} \sum_{F_1', F_2' \subset G} 1_{\{F_1' \cup F_2' \equiv H\}}. \end{aligned}$$

We now note that by definition of c_H, for each copy of H in G, there will be c_H pairs (F_1', F_2') of copies of F_1 and F_2 in G such that $F_1' \cup F_2' = H$. Therefore we can simplify the sum above to obtain:

$$\begin{aligned} X_{F_1}(G)X_{F_2}(G) &= \sum_{H \in \mathcal{H}_{F_1 F_2}} c_H \sum_{H' \subset G} 1_{\{H' \equiv H\}} \\ &= \sum_{H \in \mathcal{H}_{F_1 F_2}} c_H X_H(G), \end{aligned}$$

yielding the desired result. □

With these tools in hand, we compute the first two moments of $X_F(G)$ when $G \sim G(n, f)$.

Proposition A.1. *Fix two graphs F and F' and a random graph $G \sim G(|G|, f)$ such that $|F| + |F'| \leq |G|$. Then, we have that*

$$\begin{aligned} \mathbb{E}X_F(G) &= X_F(K_G)\mu_F(f) \\ \operatorname{cov}(X_F(G), X_{F'}(G)) &= \sum_{H \in \mathcal{H}^*_{FF'}} c_H X_H(K_G)\big(\mu_H(f) - \mu_F(f)\mu_{F'}(f)\big). \end{aligned}$$

Proof. We prove each statement in succession. To begin, call $F_1, \ldots, F_m$ the $m = X_F(K_G)$ copies of F in K_G ($m > 0$ because $|F| \leq |G|$). Then,

$$X_F(G) = \sum_{i \in [m]} 1_{\{F_i \subset G\}}.$$

Therefore, by linearity of the expectation, we have

$$\mathbb{E}X_F(G) = \sum_{i \in [m]} \mathbb{E}1_{\{F_i \subset G\}}. \tag{3}$$

Let us fix $i \in [m]$ and consider $\mathbb{E}1_{\{F_i \subset G\}}$. To do so we will use the law of total probability:

$$\begin{aligned} \mathbb{E}1_{\{F_i \subset G\}} &= \mathbb{E} \prod_{pq \in F_i} 1_{\{pq \in G\}} \\ &= \mathbb{E}\left[\mathbb{E}\left[\prod_{pq \in F_i} 1_{\{pq \in G\}} \,\middle|\, \{x_l\}_{l \in F_i}\right]\right] \\ &= \mathbb{E}\left[\prod_{pq \in F_i} f(x_p, x_q)\right] = \mu_F(f). \end{aligned}$$

Therefore, $\mathbb{E}1_{\{F_i \subset G\}}$ does not depend on i, and resuming from (3), we obtain

$$\mathbb{E}X_F(G) = \sum_{i \in [m]} \mu_F(f) = m\mu_F(f) \tag{4}$$

which is the desired result.

We now turn to the variance. Call $F'_1, \dots, F'_{m'}$ the $m' = X_{F'}(K_G)$ copies of F' in K_G ($m' > 0$ because $|F'| \leq |G|$). We first write that

$$\begin{aligned}
&\operatorname{cov}(X_F(G), X_{F'}(G)) \\
&\qquad = \mathbb{E}\left[\left(\sum_{i\in[m]} \left\{1_{\{F_i \subset G\}} - \mu_F(f)\right\}\right)\left(\sum_{i\in[m']} \left\{1_{\{F'_i \subset G\}} - \mu_{F'}(f)\right\}\right)\right] \\
&\qquad = \sum_{i\in[m], i\in[m']} \mathbb{E}\left[\big(1_{\{F_i \subset G\}} - \mu_F(f)\big)\big(1_{\{F'_{i'} \subset G\}} - \mu_{F'}(f)\big)\right]. \qquad (5)
\end{aligned}$$

Now, we observe that if F_i is disjoint form $F'_{i'}$ (i.e., $F_i \cap F'_{i'} = \emptyset$ which is possible because $|F| + |F'| \leq |G|$) then $1_{\{F_i \subset G\}}$ is independent from $1_{\{F'_{i'} \subset G\}}$ and we have that

$$\begin{aligned}
&\mathbb{E}\big[\big(1_{\{F_i \subset G\}} - \mu_F(f)\big)\big(1_{\{F'_{i'} \subset G\}} - \mu_{F'}(f)\big)\big] \\
&\qquad = \mathbb{E}\big(1_{\{F_i \subset G\}} - \mu_F(f)\big)\mathbb{E}\big(1_{\{F'_{i'} \subset G\}} - \mu_{F'}(f)\big) \\
&\qquad = \big(\mu_F(f) - \mu_F(f)\big)\big(\mu_{F'}(f) - \mu_{F'}(f)\big) \\
&\qquad = 0.
\end{aligned}$$

Therefore resuming from (5) we obtain that

$$\begin{aligned}
&\operatorname{cov}(X_F(G), X_{F'}(G)) \\
&\qquad = \sum_{\substack{i\in[m], j\in[m'] \\ F_i \cap F_{[i']'} \neq \emptyset}} \mathbb{E}\left[\big(1_{\{F_i \subset G\}} - \mu_F(f)\big)\big(1_{\{F'_{i'} \subset G\}} - \mu_{F'}(f)\big)\right] \\
&\qquad = \sum_{\substack{i\in[m], j\in[m'] \\ F_i \cap F_{[i']'} \neq \emptyset}} \left\{\mathbb{E}\left[1_{\{F_i \subset G\}} 1_{\{F'_{i'} \subset G\}}\right] - \mu_F(f)\mu_{F'}(f)\right\} \\
&\qquad = \sum_{\substack{i\in[m], j\in[m'] \\ F_i \cap F_{[i']'} \neq \emptyset}} \left\{\mathbb{E}\left[1_{\{F_i \cup F'_{i'} \subset G\}}\right] - \mu_F(f)\mu_{F'}(f)\right\}.
\end{aligned}$$

Then, by the exact same transformation we used in Lemma A.1, we know that each $F_i \cup F'_{i'}$ is in $\mathcal{H}^*_{FF'}$ (because they are not disjoint), and furthermore, for each copy of H in G there are c_H pairs F_u and F'_v such that $H = F_v \cup F'_v$, so that

$$\begin{aligned}
\operatorname{cov}(X_F(G), X_{F'}(G)) &= \sum_{H\in\mathcal{H}^*_{FF'}} c_H \sum_{H'\subset G} \left\{\mathbb{E}(1_{\{H'\equiv H\}}) - \mu_F(f)\mu_{F'}(f)\right\} \\
&= \sum_{H\in\mathcal{H}^*_{FF'}} c_H\big(\mathbb{E}X_H(G) - X_H(K_G)\mu_F(f)\mu_{F'}(f)\big) \\
&= \sum_{H\in\mathcal{H}^*_{FF'}} c_H X_H(K_G)\big(\mu_H(f) - \mu_F(f)\mu_{F'}(f)\big),
\end{aligned}$$

which is the desired result. □

We now turn to the proofs of Theorem 1 and Theorem 2. As the second generalizes the first, it is sufficient to prove the second.

Proof. We obtain the result by a joint application of the Lindeberg-Feller central limit theorem and the Cramer-Wold device. To do so, we fix $\mathbf{a} \in \mathbb{R}^{|\mathcal{F}|}$ and compute the variance of our estimator projected along $\mathbf{a}$.

Computing the variance: First recall that

$$\hat{\mu}_{\mathcal{F}}(\mathcal{G}) = N^{-1} \sum_{i \in [N]} \left(\frac{X_F(G_i)}{X_F(K_{n_i})} \right)_{F \in \mathcal{F}}.$$

Therefore, denoting "$\cdot$" the inner product and taking the expectation over $\mathcal{G}$, let

$$s_N^2 = \operatorname{var}\left(\mathbf{a} \cdot \hat{\mu}_{\mathcal{F}}(\mathcal{G})\right).$$

Then, using the independence of the G_is and the bi-linearity of the covariance, we have

$$\begin{aligned} s_N^2 &= N^{-2} \sum_{i \in [N]} \operatorname{var}\left(\mathbf{a} \cdot \left(\frac{X_F(G_i)}{X_F(K_{n_i})} \right)_{F \in \mathcal{F}} \right) \\ &= N^{-2} \sum_{i \in [N]} \sum_{F, F' \in \mathcal{F}} \frac{\mathbf{a}_F \mathbf{a}_{F'} \operatorname{cov}\left(X_F(G_i), X_{F'}(G_i)\right)}{X_F(K_{n_i}) X_{F'}(K_{n_i})}. \end{aligned}$$

To proceed, recall that for each $i \leq N$, $G_i \sim G(n_i, f_i)$. Then, we may use Proposition A.1 to obtain

$$s_N^2 = N^{-2} \sum_{i \in [N]} \sum_{F, F' \in \mathcal{F}} \frac{\mathbf{a}_F \mathbf{a}_{F'} \sum_{H \in \mathcal{H}^*_{FF'}} c_H X_H(K_{n_i}) \big(\mu_H(f_i) - \mu_F(f_i)\mu_{F'}(f_i)\big)}{X_F(K_{n_i}) X_{F'}(K_{n_i})}.$$

Because all sums are finite, we can reorder the summations, leading to

$$\begin{aligned} s_N^2 &= N^{-1} \sum_{F, F' \in \mathcal{F}} \mathbf{a}_F \mathbf{a}_{F'} \\ &\qquad \sum_{H \in \mathcal{H}^*_{FF'}} \left(N^{-1} \sum_{i \in [N]} \frac{c_H X_H(K_{n_i})}{X_F(K_{n_i}) X_{F'}(K_{n_i})} \big(\mu_H(f_i) - \mu_F(f_i)\mu_{F'}(f_i)\big) \right) \\ &= N^{-1} \sum_{F, F' \in \mathcal{F}} \mathbf{a}_F \mathbf{a}_{F'} \sum_{H \in \mathcal{H}^*_{FF'}} \omega_H(n, f; N), \end{aligned} \tag{6}$$

where

$$\omega_H(n, f; N) = N^{-1} \sum_{i \in [N]} \frac{c_H X_H(K_{n_i})}{X_F(K_{n_i}) X_{F'}(K_{n_i})} \big(\mu_H(f_i) - \mu_F(f_i)\mu_{F'}(f_i)\big).$$

Convergence of $\omega_H(n, f; N)$ To proceed we must show that $\omega_H(n, f; N)$ converges to a limit as N diverges. We will achieve this by showing that $\omega_H(n, f; N)$ is Cauchy. To do so, observe that

$$\begin{aligned}
&|\omega_H(n, f; N+1) - \omega_H(n, f; N)| \\
&\quad = \left| (N+1)^{-1} \left(N\omega_H(n, f; N) + \frac{c_H X_H(K_{n_{N+1}})}{X_F(K_{n_{N+1}}) X_{F'}(K_{n_{N+1}})} \right.\right. \\
&\qquad\qquad \left.\left. \big(\mu_H(f_{N+1}) - \mu_F(f_{N+1})\mu_{F'}(f_{N+1})\big) \right) - \omega_H(n, f; N) \right| \\
&\quad = \left| (N+1)^{-1} \frac{c_H X_H(K_{n_{N+1}})}{X_F(K_{n_{N+1}}) X_{F'}(K_{n_{N+1}})} \right. \\
&\qquad\qquad \left. \big(\mu_H(f_{N+1}) - \mu_F(f_{N+1})\mu_{F'}(f_{N+1})\big) - \frac{1}{N+1} \right|.
\end{aligned}$$

Then, we recall that $X_F(K_G) = \binom{|G|}{|F|}\mathrm{aut}(F)$ (see for instance [16]), where $\mathrm{aut}(F)$ is the number of automorphisms of F (the number of bijections from the vertex set of F to itself that preserve adjacency.) Then, as $|H| < |F| + |F'|$, we have

$$\frac{X_H(K_{n_{N+1}})}{X_F(K_{n_{N+1}}) X_{F'}(K_{n_{N+1}})} = \frac{\mathrm{aut}(H)}{\mathrm{aut}(F)\mathrm{aut}(F')} \frac{\binom{n_{N+1}}{|H|}}{\binom{n_{N+1}}{|F|}\binom{n_{N+1}}{|F'|}} \leq \frac{\mathrm{aut}(H)}{\mathrm{aut}(F)\mathrm{aut}(F')}.$$

As furthermore f_{N+1} is bounded by 1, we have

$$\left|\big(\mu_H(f_{N+1}) - \mu_F(f_{N+1})\mu_{F'}(f_{N+1})\big)\right| \leq 1,$$

leading to $|\omega_H(n, f; N+1) - \omega_H(n, f, N)| \leq \frac{C}{N+1}$ for a constant C equal to $1 + c_H \mathrm{aut}(H)/\mathrm{aut}(F)\mathrm{aut}(F')$. Thus, the sequence $\omega_H(n, f; N)$ is Cauchy, and we may call $\omega_H(n, f)$ its limit; i.e.,

$$\lim_{N\to\infty} \omega_H(n, f; N) = \omega_H(n, f).$$

Then, resuming from (6), and writing $\Sigma_{\mathcal{F}}$ the matrix indexed by $\mathcal{F}$ such that

$$(\Sigma_{\mathcal{F}}(n, f))_{FF'} = \sum_{H \in \mathcal{H}^*_{FF'}} \omega_H(n, f),$$

we have

$$\lim_{N\to\infty} N s_N^2 = \mathbf{a}^\top \Sigma_{\mathcal{F}}(n, f)\mathbf{a}.$$

Satisfying the Lindeberg-Feller condition To invoke the Lindeberg-Feller central limit theorem, we must show that our sequence verifies the so called Lindeberg-Feller condition. Recall that the sequence under study is

$$Y_i := \mathbf{a} \cdot (X_F(G_i)/X_F(K_{n_i}))_{F\in\mathcal{F}} - \mathbf{a} \cdot \mu_{\mathcal{F}}(f_i),$$

and that the variance of the partial sum is $N^2 s_N^2$. Therefore, the Lindeberg-Feller condition we need to satisfy is the following:

$$\forall \epsilon > 0 \lim_{N\to\infty} \frac{1}{N^2 s_N^2} \sum_{i\in[N]} \mathbb{E}\left[Y_i 1_{\{|Y_i|>\epsilon N s_N\}}\right] = 0. \tag{7}$$

To verify the condition we first fix $\epsilon > 0$. Then, observe that since for each i and F we have $X_F(G_i)/X_F(K_{n_i}) \leq 1$ and $\mu_F(f_i) \leq 1$, we have that $|Y_i| \leq 2\|\mathbf{a}\|_1$ by the triangle inequality. Therefore, as $N s_N \to \infty$ as N grows, we may fix an N such that for all $N' > N$ we have $\|\mathbf{a}\|_1 \leq \epsilon N s_N$. In this setting, the sum in (7) is equal to zero for all $N' > N$, and the condition is verified. Therefore, we have that

$$\frac{1}{N s_N} \sum_i Y_i \to \text{Normal}(0, 1).$$

To conclude, and reverting to the notation of the statement of the Theorem, we have that for any $\mathbf{a}$

$$\sqrt{N}\big(\mathbf{a} \cdot \hat{\mu}_{\mathcal{F}}(\mathcal{G}) - \mathbf{a} \cdot \mu_{\mathcal{F}}(f)\big) \to \text{Normal}(0, \mathbf{a}^\top \mathbf{\Sigma}_{\mathcal{F}}(n, f)\mathbf{a}),$$

which is sufficient to obtain the claimed result for the limit in distribution. □

Before we proceed, we consider a simulation experiment to determine how large the sample size N must be for the asymptotic limit to be a satisfactory approximation of the statistic's distribution. We present our result in Fig. 5. There, we observe that fairly small N, on the order of 100 even with small n_i-s, may be sufficient. Furthermore, results on large graphs [4] suggest that using larger networks would make this convergence even faster.

We now turn to the proof of Corollary 1.

Proof. The result follows almost immediately from Theorem 1 and Slutsky's theorem.

First, since $|\mathcal{G}|$ and $|\mathcal{G}'|$ tend to infinity and both n and n' are large enough, we have that

$$\begin{cases} \sqrt{|\mathcal{G}|}\big(\hat{\mu}_{\mathcal{F}}(\mathcal{G}) - \mu_{\mathcal{F}}(f)\big) \to \text{Normal}\big(0, \Sigma_{\mathcal{F}}(n, f)\big), \\ \sqrt{|\mathcal{G}'|}\big(\hat{\mu}_{\mathcal{F}}(\mathcal{G}') - \mu_{\mathcal{F}}(f)\big) \to \text{Normal}\big(0, \Sigma_{\mathcal{F}}(n', f)\big). \end{cases}$$

Then, as both samples are independent, any linear combinations will still be multivariate Gaussian. Therefore, we multiply the first line by $\sqrt{|\mathcal{G}'|}/\sqrt{|\mathcal{G}| + |\mathcal{G}'|}$, the second by $\sqrt{|\mathcal{G}|}/\sqrt{|\mathcal{G}| + |\mathcal{G}'|}$ (as both ratios are in $(0, 1)$ the limit in distribution is unaffected), and take their difference, to obtain

$$\begin{aligned} &\frac{\sqrt{|\mathcal{G}||\mathcal{G}'|}}{\sqrt{|\mathcal{G}| + |\mathcal{G}'|}}\big(\hat{\mu}_{\mathcal{F}}(\mathcal{G}) - \hat{\mu}_{\mathcal{F}}(\mathcal{G})\big) \\ &\qquad \to \text{Normal}\left(0, \lim_{|\mathcal{G}|,|\mathcal{G}'|\to\infty} \left\{ \frac{|\mathcal{G}'|}{|\mathcal{G}| + |\mathcal{G}'|} \Sigma_{\mathcal{F}}(n, f) + \frac{|\mathcal{G}|}{|\mathcal{G}| + |\mathcal{G}'|} \Sigma_{\mathcal{F}}(n', f) \right\}\right) \end{aligned}$$

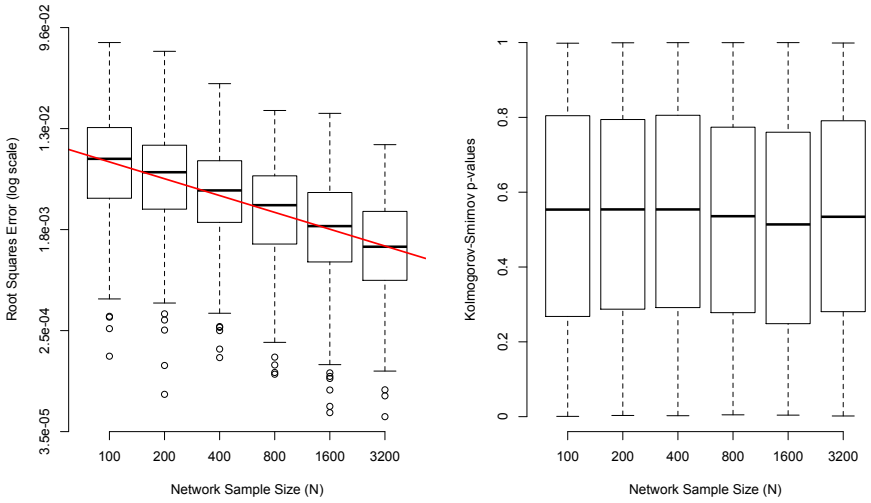

Figure 5: Assessing the quality of the asymptotic approximation. We sample 500 replicates of network samples of sizes N ranging from 100 to 3200, each of which is drawn from a random 2-block blockmodel and the sizes of the networks are fixed by the digits of π plus 8. On each sample, we evaluate the procedure presented in Fig. 2. We first consider how the average mean squared error to the mean shrinks as N increases (left plot, trend in red solid line). There, we observe rate of convergence in line with our theoretical results, and with many outliers achieving much better convergence. Then, we consider whether the Kolmogorov-Smirnov test rejects the null of the p-values across samples are uniformly distributed. There we observe that already for samples of size 100, we fail to reject the null of a uniform distribution.

$$\to \text{Normal}\big(0, \Sigma_{\mathcal{F}}(n, n', f)\big),$$

which is the desired limit in distribution.

To obtain the consistent estimator of $\Sigma_{\mathcal{F}}(n, n', f)$ we first recall that for $F, F' \in \mathcal{F}$

$$\big(\Sigma_{\mathcal{F}}(n, f)\big)_{FF'} = \sum_{H \in \mathcal{H}^*_{FF'}} \omega_H(n)\big(\mu_H(f) - \mu_F(f)\mu_{F'}(f)\big).$$

Then, with

$$\omega_H(n, n') = \lim_{|\mathcal{G}|, |\mathcal{G}'| \to \infty} \left(\frac{|\mathcal{G}'|}{|\mathcal{G}| + |\mathcal{G}'|} \omega_H(n) + \frac{|\mathcal{G}|}{|\mathcal{G}| + |\mathcal{G}'|} \omega_H(n') \right),$$

we have

$$\big(\Sigma_{\mathcal{F}}(n, n', f)\big)_{FF'} = \sum_{H \in \mathcal{H}^*_{FF'}} \omega_H(n, n')\big(\mu_H(f) - \mu_F(f)\mu_{F'}(f)\big).$$

There observe that:

– As $|\mathcal{G}|$ and $|\mathcal{G}'|$ grow we have that

$$\begin{aligned} \omega_H(n, n'; |\mathcal{G}|, |\mathcal{G}'|) &= \frac{|\mathcal{G}'|}{|\mathcal{G}| + |\mathcal{G}'|} \omega_H(n; |\mathcal{G}|) + \frac{|\mathcal{G}|}{|\mathcal{G}| + |\mathcal{G}'|} \omega_H(n', |\mathcal{G}'|) \\ &\to \omega_H(n, n'). \end{aligned}$$

– All the $\mu_H(f)$ in $H \in \mathcal{H}_{FF'}$ may be estimated by

$$\hat{\mu}_H(\mathcal{G} \cup \mathcal{G}') = \frac{1}{|\mathcal{G} \cup \mathcal{G}'|} \sum_{G \in \mathcal{G} \cup \mathcal{G}'} \frac{X_H(G)}{X_H(K_G)}.$$

Furthermore, both $\mu_F(f)$ and $\mu_{F'}(f)$ may be estimated in the same way.

Then, by a direct application of the Slutsky's theorem, we have that

$$\big(\hat{\Sigma}_{\mathcal{F}}(n, n', f)\big)_{FF'} = \sum_{H \in \mathcal{H}^*_{FF'}} \omega_H(n, n'; |\mathcal{G}|, |\mathcal{G}'|)\big(\hat{\mu}_H(\mathcal{G} \cup \mathcal{G}') - \hat{\mu}_F(\mathcal{G} \cup \mathcal{G}')\hat{\mu}_{F'}(\mathcal{G} \cup \mathcal{G}')\big)$$

is an asymptotically normal estimator of $\big(\Sigma_{\mathcal{F}}(n, n', f)\big)_{FF'}$ converging at rate $\sqrt{|\mathcal{G}| + |\mathcal{G}'|}$, which yields the claimed result. □

APPENDIX B: Methods to produce Figs. 2–5

The code to produce all figures and tables is available from the authors. All computation are done in `R`, and rely on the packages `igraph` and `mixtools` [25–27]. To ease presentation we always denote K the number of blocks in the blockmodel, or flexible blockmodel, under consideration. Furthermore, we call $\pi[k]$ the k-th digit of π.

Subgraph counting To count the number of copies of , , , we use the formulas presented in [19]. Specifically, with $A^{(k)}_{ij}$ the i,j-th entry of the k-th power of the adjacency matrix of G, we use

$$\begin{cases} X_{\bullet\!-\!\bullet}(G) &= \frac{1}{2}\sum_{i,j\in[|G|]} A_{ij} \\ X_{\triangle}(G) &= \frac{1}{6}\sum_{i\in[|G|]} A^{(3)}_{ii} \\ X_{\square}(G) &= \frac{1}{8}\left(\sum_{i,j\in[|G|]} A^{(3)}_{ij}A_{ij} - 4\sum_{i\in[|G|]}\binom{A^{(2)}_{ij}}{2} - 2X_{\bullet\!-\!\bullet}(G)\right). \end{cases}$$

Fig. 2 We first describe the network sample. The sample is such that for each $i \leq 300$, $G_i \sim G(\pi[i] + 40, f_B)$, where

$$f_B : (0,1)^2 \to [0,1]; \forall u, v, \quad f_B(u,v) = B_{1_{\{4u\geq 1\}}+1, 1_{\{4v\geq 1\}}+1},$$

and

$$B = \begin{pmatrix} .06 & .02 \\ .02 & .10 \end{pmatrix}.$$

We now describe the construction of the confidence region and p-value under the null. The first task is to compute the subgraph densities (the $\mu_F(f_B)$ and $\mu_H(f_B)$). To do so, we use the formulas presented in [4,28] to compute subgraph densities in the blockmodel. More specifically, they show that for a graphon blockmodel f with K blocks such that the probability of being in each block is π_i, and the block matrix is B, then

$$\mu_F(f) = \sum_{i_i,\ldots,i_{|F|}\in[K]} \prod_{j\in[|F|]} \pi_j \prod_{pq\in F} B_{i_p i_q}.$$

Practically, this is achieved by imbricated loops.

That we can compute the subgraph densities under the null directly implies that we can compute both $\mu_{\mathcal{F}}(f_B)$ and $\Sigma_{\mathcal{F}}(f_B)$. Then, the construction of the confidence ellipse as well as of the p-value using the Mahalanobis distance are classical statistical inference methods which we need not describe here.

The blue region is built with $f_{B'}$ where

$$f_{B'} : (0,1)^2 \to [0,1]; \forall u, v, \quad f_{B'}(u,v) = B'_{1_{\{4u\geq 1\}}+1, 1_{\{4v\geq 1\}}+1},$$

and

$$B' = \begin{pmatrix} .06 & .03 \\ .03 & .10 \end{pmatrix}.$$

Fig. 3 We first describe the network sample. The sample is such that for each $i \leq 200$, $G_i \sim G(\pi[k] + 30, f_B)$, where

$$f_B : (0,1)^2 \to [0,1]; \forall u, v, \quad f_B(u,v) = B_{\lfloor 3u\rfloor+1, \lfloor 3v\rfloor+1},$$

and

$$B = \begin{pmatrix} .06 & .02 & .02 \\ .02 & .06 & .02 \\ .02 & .02 & .06 \end{pmatrix}.$$

We now describe how we produced the embedding shape and its confidence region. First, to produce the embedding shape, we compute $\mu_{\mathcal{F}}(f)$ for a a large number of elements of $f \in \mathcal{D}(B)$. Practically, we parametrize $\mathcal{D}(B)$ by a K-dimensional vector π, the entries of which are the sizes of each block, and call the associated graphon f_π. Then, we build a grid S over the K dimensional simplex of step-size .01, and for each $\pi \in S$ we compute $\mu_{\mathcal{F}}(f_\pi)$. To produce the confidence region, we produce the confidence ellipse for each $\pi \in S$ which we achieve by computing $\Sigma_{\mathcal{F}}(n, f_\pi)$ for each π in our grid S.

The blue region is built with $f_{B'}$ where

$$f_{B'} : (0,1)^2 \to [0,1]; \forall u, v, \quad f_{B'}(u,v) = B'_{\lfloor 3u \rfloor + 1, \lfloor 3v \rfloor + 1},$$

and

$$B' = \begin{pmatrix} .05 & .04 & .04 \\ .04 & .05 & .04 \\ .04 & .04 & .05 \end{pmatrix}.$$

Fig. 4 We first describe the network sample. The sample is such that independently for each $i \leq 1000$, $G_i \sim G(\pi[i] + 50, f_B)$, where f_B is drawn uniformly at random over $\mathcal{D}(B)$ with

$$B = \begin{pmatrix} .7 & .2 & .2 \\ .2 & .7 & .2 \\ .2 & .2 & .7 \end{pmatrix}.$$

We now describe how we produced the surface where $\mu_{\mathcal{F}}(f)$ may live as well as its confidence region. First, we know that $\mu_{\mathcal{F}}(f)$ may realize any point in the convex hull of the embedding shape. Therefore, we build the embedding shape as for Fig. 3, and then present its convex hull. Building the confidence region is achieved using the following reasoning, assuming that the diagonal of B dominates the off diagonal entries ($\max_{i \in [K]} B_{ii} \geq \max_{i \neq j} B_{ij}$):

- Observe that for any graphon f and $H \in \mathcal{H}_{FF'}$, $\mu_H(f) \geq \mu_F(f)\mu_{F'}(f)$. This can for instance be directly recovered from the formulas in Lemma A.1. Then, for any sequences n and f, we have that $(\Sigma_{\mathcal{F}}(n, f))_{FF'} \geq 0$, so that all correlations are positive, and the greatest amplitude of the confidence ellipse will be found on its first quadrant; i.e., the variance of $a \cdot \mu_{\mathcal{F}}(\mathcal{G})$ will be maximal for some $a \in \mathbb{R}^{|\mathcal{F}|}$ such that $a_F > 0$ for all $F \in \mathcal{F}$.
- For any sequences n and f, we have that

$$\begin{aligned}
(\Sigma_{\mathcal{F}}(n, f))_{FF'} &= \sum_{F,F' \in \mathcal{F}} \sum_{i \in [N]} \frac{c_H X_H(K_{n_i})}{X_F(K_{n_i}) X_{F'}(K_{n_i})} \big(\mu_H(f_i) - \mu_F(f_i)\mu_{F'}(f_i)\big) \\
&\leq \sum_{F,F' \in \mathcal{F}} \sum_{i \in [N]} \frac{c_H X_H(K_{n_i})}{X_F(K_{n_i}) X_{F'}(K_{n_i})} \mu_H(f_i) \\
&\leq \sum_{F,F' \in \mathcal{F}} \sum_{i \in [N]} \frac{c_H X_H(K_{n_i})}{X_F(K_{n_i}) X_{F'}(K_{n_i})} \max_{f \in \mathcal{D}(B)} \mu_H(f) := (\Sigma^{\dagger}_{\mathcal{F}}(n, f))_{FF'}.
\end{aligned}$$

Then, for any a in the first quadrant we have that

$$a^\top \Sigma_{\mathcal{F}}(n,f)a \leq a^\top \Sigma^\dagger_{\mathcal{F}}(n,f)a.$$

Therefore, the maximal radius of the ellipse associated with $\Sigma^\dagger_{\mathcal{F}}(n,f)$ on the first quadrant is larger than that induced by $\Sigma_{\mathcal{F}}(n,f)$. Thus, by the previous item, the maximal radius of $\Sigma^*_{\mathcal{F}}(n,f)$ is larger than that of $\Sigma_{\mathcal{F}}(n,f)$. It follows that the auxiliary sphere of the ellipse associated with $\Sigma^\dagger_{\mathcal{F}}(n,f)$ contains that associated with $\Sigma_{\mathcal{F}}(n,f)$.

– Finally, under our assumption of diagonal dominance in B, we have that $\max_{f\in\mathcal{D}(B)} \mu_H(f)$ is realized by the graphon such that all nodes are in the same block r with the highest probability of within-block connection; i.e., as parametrized by π, the maximum would be realized by $\pi^* = (1_{\{i=r\}})_{i\in[K]}$ where r is such that $B_{rr} = \max_{i\in[K]} B_{ii}$. In this final setting, we need only compute the $\mu_H(f_{\pi^*})$, where f_{π^*} describes an Erdős-Rényi random graph. This makes the computation straightforward as we then have $\mu_H(f_{\pi^*}) = B^k_{rr}$, where k is the number of edges in H.

Following this argument, our confidence region is the union of the auxiliary spheres of the ellipse associated with $\Sigma^*_{\mathcal{F}}(n,f)$ centered at each point in the convex hull of the embedding shape.

Finally, the blue region is built with

$$B = \begin{pmatrix} .8 & .5 & .5 & .5 \\ .5 & .8 & .5 & .5 \\ .5 & .5 & .8 & .5 \\ .5 & .5 & .5 & .8 \end{pmatrix}.$$

References

[1] W. R Gray, J. A. Bogovic, J. T Vogelstein, B. A. Landman, J. L. Prince, and R. J. Vogelstein. Magnetic resonance connectome automated pipeline: an overview. *IEEE pulse*, 3(2):42–48, 2012.

[2] P. D. Hoff, A. E. Raftery, and M. S. Handcock. Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.*, 97(460):1090–1098, 2012.

[3] Q. Ho, A. P. Parikh, and E. P. Xing. A multiscale community blockmodel for network exploration. *J. Amer. Statist. Assoc.*, 107(499):916–934, 2012.

[4] P. J. Bickel, A. Chen, and E. Levina. The method of moments and degree distributions for network models. *Ann. Statist.*, 39:2280–2301, 2012.

[5] D. L. Sussman, M. Tang, D. E. Fishkind, and C. E. Priebe. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *J. Amer. Statist. Assoc.*, 107:1119–1128, 2012.

[6] S. Bhattacharyya and P. J. Bickel. Subsampling bootstrap of count features of networks. *Ann. Statist.*, 43:2384–2411, 2015.

[7] M. Tang, D. L. Sussman, and C. E. Priebe. Universally consistent vertex classification for latent positions graphs. *Ann. Statist.*, 41:1406–1430, 2013.

[8] S. C. Olhede and P. J. Wolfe. Network histograms and universality of blockmodel approximation. *Proc. Natl. Acad. Sci. USA*, 111:14722–14727, 2014.

[9] B.K. Fosdick and P. D. Hoff. Testing and modeling dependencies between a network and nodal attributes. *J. Amer. Statist. Assoc.*, 110(511):1047–1056, 2015.

[10] O. Klopp, A. B. Tsybakov, and N. Verzelen. Oracle inequalities for network models and sparse graphon estimation. *Ann. Statist.*, Forthcoming, 2016.

[11] D. Asta and C. R. Shalizi. Geometric network comparison. arXiv:1411.1350, 2014.

[12] M. Tang, A. Athreya, D. L. Sussman, V. Lyzinski, and C. E. Priebe. A nonparametric two-sample hypothesis testing for random dot product graphs. *Bernoulli*, Forthcoming, 2017.

[13] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: Simple building blocks of complex networks. *Science*, 298:824–827, 2002.

[14] T. Ali, T. Rito, G. Reinert, F. Sun, and C. M. Deane. Alignment-free protein interaction network comparison. *Bioinformatics*, 30:i430–i437, 2014.

[15] P. Erdös and A. Rényi. On random graphs I. *Publicat. Mathemat.*, 6:290–297, 1959.

[16] B. Bollobás and O. Riordan. Metrics for sparse graphs. In S. Huczynska, J. D. Mitchell, and C. M. Roney-Dougal, editors, *Surveys in Combinatorics 2009*, pages 211–287. Cambridge University Press, Cambridge, UK, 2009.

[17] L. Lovász. *Large Networks and Graph Limits.* Am. Mathematical Soc., Providence, RI, 2012.

[18] A. Rucinski. When are small subgraphs of a random graph normally distributed? *Probab. Theory Related Fields*, 78(1):1–10, 1988.

[19] N. Alon, R. Yuster, and U. Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997.

[20] T. Hočevar and J. Demšar. A combinatorial approach to graphlet counting. *Bioinformatics*, 30:559–565, 2014.

[21] M. Jha, C. Seshadhri, and A. Pinar. Path sampling: A fast and provable method for estimating 4-vertex subgraph counts. In *Proceedings of the 24th International Conference on World Wide Web*, WWW '15, pages 495–505, New York, NY, USA, 2015. ACM.

[22] J. T. Vogelstein, R. J. Vogelstein, and C. E. Priebe. Are mental properties supervenient on brain properties? *Scientific Reports*, 1, 2011.

[23] D. Koutra, J. T. Vogelstein, and C. Faloutsos. DELTACON: A principled massive-graph similarity function. In *Proceedings of the SIAM International Conference in Data Mining. Society for Industrial and Applied Mathematics*, pages 162–170. SIAM, 2013.

[24] A. Pinar, C. Seshadhri, and V. Vishal. Escape: Efficiently counting all 5-vertex subgraphs. arXiv:1610.09411, 2016.

[25] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016.

[26] G. Csardi and T. Nepusz. The igraph software package for complex network research. *InterJournal*, Complex Systems:1695, 2006.

[27] T Benaglia, D. Chauveau, D. R. Hunter, and D. Young. mixtools: An R package for analyzing finite mixture models. *Journal of Statistical Software*, 32(6):1–29, 2009.

[28] M. Coulson, R. E. Gaunt, and G. Reinert. Poisson approximation of subgraph counts in stochastic block models and a graphon model. arXiv:1509.07754, 2016.